



EvalFest How-Tos for Conducting an Independent Samples T-test using R and the RStudio Editor

Independent Samples T-test

1. Intro and Data Setup
 - a. If you've watched our video that explains the concept of an independent samples t-test, you may recall that an independent-samples t-test helps you figure out when you really do have apples and oranges in your data where apples and oranges are different groups of people. Take a look at the file containing the code we will use to tell Rstudio that we want to run an independent samples t-test. If you've watched our video that shows you how to set up the data and variables in R, you will notice that lines 1-3, 6, and 7 are part of the initial setup.
2. Running the Independent Samples T-test
 - a. In this example, we are going to take a look at average Overall Rating of the Festival by Gender, using two groups for gender – female and male. The Overall Rating question asked festival attendees, “How would you rate today's event overall?” Festival goers could choose an answer on a 5-point scale from Poor to Excellent.
 - b. Here, Gender is the predictor variable – we want to know if Gender predicts differences in Overall Rating. Overall Rating is our outcome variable – it is the variable that we want to learn more about.
 - c. In this file, Line 10 shows the code that conducts the independent samples t-test.
 - d. Notice that we have our alternative hypothesis defined as ‘two.sided’. This means that we don't know if Males or Females will have a higher average Overall Rating. This is the default in R and is not necessary if you want to conduct a two-sided test. It is also the default that you will use unless you have a theoretical or historical reason to believe that you know which group will have the higher rating.
 - i. If you want to conduct an independent samples t-test in which you indicate that the mean for group 1 is greater than the mean for group 2 (in this scenario, your alternate hypothesis would state that the mean overall rating for females is significantly greater than the mean overall rating for males), you can replace “two.sided” with “greater”.
 - ii. If you want to conduct an independent samples t-test in which you indicate that the mean for group 1 is less than the mean for group 2 (in this scenario, your alternate hypothesis would state that the mean overall rating for females is significantly less than the mean overall rating for males), you can replace “two.sided” with “less”.

- iii. R assumes that the first variable you list in the `t.test` line of code contains the data for group 1 – your apples group, and the 2nd variable contains the data for group 2 – your oranges.
 - e. Also notice that we have included the statement `paired=FALSE`. This tells R that we are conducting an independent samples t-test. This is the default in R and is not necessary if you want to conduct an independent samples t-test.
 - i. There is also something called a paired-samples t-test, and you would want to change this from `FALSE` to `TRUE` if that is the test you wanted to conduct.
- 3. Interpretation of the Independent Samples T-test Output
 - a. Running line 10 gives you output – these are your results. RStudio provides the value of *t*, the degrees of freedom, and the *p*-value, among other things. The t-test uses the average rating for each group, and the variation in those ratings, to figure out if you have apples and oranges in your data, or if you really just have different kinds of apples. If the *p*-value is less than .05 then that means you DO have differences in your groups – apples and oranges. In more statistical terms, you use the *p*-value to determine if the mean of the outcome variable is significantly different between the groups in the predictor variable. For this scenario, we can see that our *p* value is 5.199×10^{-8} , or .00000005199. That number is less than .05 and so that means that we have a significant difference in the average ratings provided by female and male attendees.
 - b. Note that the *p*-value simply tells you if the ratings are different – it doesn't tell you which group gave the higher rating. To figure that out, we have to go back to look at the mean ratings for each group. We see that the mean overall rating for females is 4.44, which we found is significantly higher than the mean overall rating of 4.33 for males.
 - i. So now what? How do you think about interpreting these results? Remember that these ratings were made on a five-point scale, where 1 was Poor and 5 was Excellent. Both of our average ratings were above 4.0, and so somewhere between Very Good and Excellent on the scale. When interpreting these kinds of results, you want to think about the kinds of ratings you have overall and the group differences. So, for these results, we might say something like: Ratings for the Festival were quite positive. There was a significant difference in how males and females rated the event. Males rated the Festival as very good, on average, with a rating of 4.33. Female ratings were significantly higher, and thus even closer to the top of the scale at 4.44.

